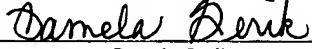


**PATENT  
5201-27600  
03-1237**

CERTIFICATE OF EXPRESS MAIL UNDER 37 C.F.R. § 1.10
"Express Mail" mailing label number <u>EV403685495US</u>
DATE OF DEPOSIT <u>December 5, 2003</u>
I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" Service Under 37 C.F.R. §1.10 on the date indicated above and is addressed to: Commissioner for Patents and Trademarks, <b>BOX PATENT APPLICATION</b> , Washington, D.C. 20231
 Pamela Gerik

**LOW POWER MEMORY CONTROLLER WITH LEADED DOUBLE DATA  
RATE DRAM PACKAGE ARRANGED ON A TWO LAYER PRINTED CIRCUIT BOARD**

By:

Eric Hung  
470 Navaro Way #106  
San Jose, California 95134  
Citizenship: U.S.

Geeta K. Desai  
19700 Farwell Avenue  
Saratoga, California 95070  
Citizenship: India

Vijendra Kuroodi  
20875 Valley Green Drive #91  
Cupertino, California 95014  
Citizenship: India

## **BACKGROUND OF THE INVENTION**

### **1. Field of Invention**

5           This invention relates to semiconductor memory and, more particularly, to a low power memory controller placed on a two-layer printed circuit board (PCB) with a double data rate (DDR) DRAM, both of which are packaged with leads extending outward from the package and secured to pads on the PCB for a lower cost semiconductor memory system.

### **2. Description of Related Art**

          The following descriptions and examples are not admitted to be prior art by virtue of their inclusion within this section.

15           There are numerous types of memories available in the marketplace. For example, large volumes of data can be stored in magnetic memory, such as a hard disk drive. Lesser quantities of data can be stored in memory arranged upon an integrated circuit, oftentimes referred to as “semiconductor memory.” Semiconductor memory is typically arranged closer to the core logic unit, or execution unit, than the hard disk drive and can be accessed much faster than the disk drive.

          Common to semiconductor memory is an array of storage cells. Depending on the function of the semiconductor memory, each storage cell can have a significantly different architecture and function. For example, semiconductor memory can be volatile or non-volatile. Types of volatile memory include memory which must be periodically refreshed (DRAMs) or memory that will lose its programmed state if power is removed (SRAMs).

The differences between SRAMs and DRAMs are fairly significant. For example, each cell of SRAM includes latch and pass transistors. Conversely, each cell of DRAM involves simply one transistor. While DRAMs are denser than SRAMs, DRAMs require additional support circuitry to coordinate the access of each cell, along with the need to periodically refresh that cell. The faster access time of SRAMs allow their as the primary cache of the execution unit, whereas DRAMs are generally used as the main semiconductor memory. SRAMs are, therefore, placed on the same address and data bus as the execution unit, whereas DRAMs are linked to the execution unit by a memory controller and stand-alone memory bus.

10

In addition to having a memory controller between the array of storage cells (DRAMs) and the execution unit, the DRAMs themselves take on numerous forms: fast page DRAMs, extended data out DRAMs, burst extended data out DRAMs, and the more recent synchronous DRAMs (SDRAMs). SDRAMs take advantage of the fact that memory accesses by the execution unit are typically sequential. SDRAMs are thereby designed to fetch all bits within a particular burst in sequential fashion by allowing the column address to be incremented sequentially and in sync with the system clock of the execution unit or processor.

15

With the increased access speed of the SDRAM came yet another enhancement to DRAMs in general. Instead of providing source-synchronous data capture at the clock frequency, double data rate (DDR) DRAMs allow data to be captured at twice the rate of the clock frequency. Instead of capturing data once per cycle as in single data rate (SDR) DRAMs, DDR DRAMs utilize a 2n prefetch architecture to allow data to be captured twice per clock cycle. Details of the difference between SDR and DDR are set forth in "General DDR SDRAM Functionality," Micron Technology 2001 (herein incorporated by reference).

20

25

While both SDR and DDR DRAMs include the same core memory array of cells, the input/output (I/O) interface is considerably different. For example, DDR utilizes a true and complementary pair of clock signals to formulate the triggering rising and falling

30

clock edges, and also utilizes a data strobe signal (DQS) to drive the data signal (DQ) to and from the memory banks.

One benefit of DDR is the higher speed at which data can be read from or written to the storage cells. Unfortunately, however, the increases speed of the I/O interface within the memory controller limits the available selection of available packing techniques. For example, when the data transfer rate exceeds several 100 MHz, relatively high inductance occur on the data signals sent to and from the memory controller. In an attempt to reduce the inductance, the data pins are kept as short as possible and are placed as close to the ground plane as possible. By shortening the data pins or leads extending from the controller and/or DRAMs and placing those pins near the ground plane, the current loop area and the magnetic flux are minimized. As described in "Circuits, Interconnections, and Packaging for VLSI," 1990, pg. 321 (herein incorporated by reference), packages which incur the poorest lead inductance are dual in-line packages (DIPs) or small outline packages (SOPs) which have approximately 3-50 nH inductance. Packages that have the smallest lead inductance (.025-1 nH) are those that employ flip-chip technology.

To take full advantage of the higher transfer speeds of DDR, the integrated circuit having the memory controller also has preferably an array of solder bumps arranged on the upper surface of the integrated circuit. The integrated circuit is then flipped over so that the solder bumps are in alignment with solder balls on a single or multi-layer substrate. The "flipped" integrated circuit is then heated along with the substrate, to establish a reflow bond between the integrated circuit and the substrate. Flip-chip mounting with collapsible solder balls provides the least parasitic inductance and capacitance, generally less than 1 nH and less than 1 pF. See "Circuits, Interconnections, and Packaging for VLSI," 1990, pg. 91 (herein incorporated by reference).

While flip-chip mounting techniques enjoy lessened parasitic inductance and capacitance, flip-chip mounting is nonetheless relatively expensive. First, the integrated circuit must be laid out with bonding pads on an upper surface, and one or more ground

planes either on the integrated circuit near the solder bumps or on the substrate near the solder balls. Second, the added multi-layer substrate and the added encapsulate process needed to secure the collapsible solder balls increases the overall complexity of the packaged integrated circuit. Third, in addition to the specially configured array of solder bumps and the encapsulate between the collapsible solder balls and solder bumps, flip-chip mounting also requires additional mounting of the substrate-connected integrated circuit to a much larger printed circuit board (PCB).

Instead of requiring a special purpose multi-layer substrate, and flip-chip arrangement of solder balls and bumps, it would be desirable to utilize a less expensive packaging technique. For example, use of a lead frame and wire bonds to form a packaged integrated circuit, and simply mounting the leads onto a PCB without any intervening special-purpose substrate would not only be faster, but less costly to implement than flip-chip mounting. However, if conventional lead frames and wire bonding packages are to be used, it would be desirable to implement modifications to the overall data rate transfer of a DDR DRAM.

In addition to reducing the overall packaging costs of the integrated circuit, additional cost savings must be implemented on the PCB to further reduce costs of the interconnect solution. For example, it would be desirable to introduce DDR transfers in a lower cost small outline package (SOP), with a minimum number of layers formed on the PCB. The desired solution would achieve the benefits of DDR, yet also reduce the costs of packaging, interconnect, and PCB complexity that currently exists in the flip-chip mounting solution.

## SUMMARY OF THE INVENTION

The problems outlined above are in large part solved by implementing a memory module that comprises a DDR DRAM, a memory controller, and an execution unit  
5 mounted on a PCB. The execution unit and memory controller can be formed on the same integrated circuit, and packaged in a SOP with the integrated circuit mounted on a lead frame paddle, and the leads forming a part of the lead frame itself. The leads will extend from the packaged integrated circuit, and can be solder-bonded directly to the PCB using, for example, various surface mount techniques. Bonding pads on the  
10 integrated circuit are preferably coupled to bond fingers on the lead frame via, for example, wires. In addition to the packaged execution unit and memory controller mounted to the PCB, various DRAMs can also be packaged and mounted to the PCB.

The PCB is preferably a multi-layer PCB with a maximum of two conductive  
15 layers or planes. The conductive layers are formed by selectively removing a thin sheet of conductive material to produce a plurality of coplanar trace conductors. The trace conductors on each layer comprise conductors that receive signals, a power supply, and ground. The power supply, signal, and ground conductors on each layer are separated by a dielectric layer, and preferably one of the pair of conductive layers is on an outer  
20 surface of the PCB. The exposed conductive layer may have one or more conductive pads that receive corresponding leads of the SOP integrated circuit.

By utilizing a lead frame and wire bonds to form the packaged integrated circuits, substantial cost savings can be realized. Moreover, using only two conductive layers  
25 rather than four or more can realize additional cost savings. The cost savings, however, are mitigated if the added parasitic inductance and capacitance are unduly large. In order to reduce the parasitic inductance and capacitance on the leads of the SOP surface mounted to the bonding pads on the upper conductive layer of the PCB, a clock throttling technique is used herein to control the data transfer rate.

30

Clock throttling is used with the advantages of DDR DRAMs. Currently, the marketplace has demonstrated significant demand for DDR as opposed to SDR. As the demand for DDR asynchronous DRAM or DDR SDRAM continues to increase, the costs of DDR will be less than SDR asynchronous DRAM or SDRAM. While further cost savings can be achieved using DDR, clock throttling of DDR is preferably used so that the data transfer rate and the latch speed of the I/O interface of the memory controller does not become unduly fast so as to produce undesirably high parasitic inductance and capacitance. The present memory module is thereby a cost benefit solution to conventional DDR DRAMs that utilize more expensive flip-chip mounting technologies or more expensive SDR DRAMs currently available in the marketplace.

According to one embodiment, an integrated is provided. The integrated circuit includes an execution engine, memory controller, or both an execution unit and a memory controller. The execution engine is preferably clocked at a first clock rate, and the memory controller is preferably clocked at a second clock rate. The second clock rate is desirably less than the first clock rate. A plurality of pins are configured on the integrated circuit and, preferably, extend from the integrated circuit as part of a lead frame. The pins are adapted to transfer data to and from the memory controller on both the rising and falling edges of the second clock signal transitioning at the second clock rate.

A multiplexer can be used to throttle the memory controller to capture data on the plurality of pins at a lessened second clock rate. The multiplexer can be configured to select the second clock rate rather than the first clock rate when, for example, the memory controller receives a power supply voltage. In addition to the memory controller being clocked at a lessened, second clock rate for capturing data on the rising and falling edges of the second clock signal, the memory controller is also preferably coupled to a power supply voltage less than 2.5 volts, for example. The memory controller can preferably utilize differential I/O signals such that the I/O interface of the memory controller receives a power supply in accordance with the stub-series terminated logic

(SSTL\_2) standard, and not the higher voltage levels (low voltage TTL, or “LVTTL”) associated with SDR.

According to another embodiment, an electronic system is provided. The system includes a substrate (herein referred to as a PCB) that comprises a maximum of two conductive layers spaced from each other by a dielectric layer. An execution engine clocked at a first clock rate is connected to possibly a first conductor that extends partially across one of the two conductive layers. A memory controller clocked at a second clock rate can be coupled to a second conductor extending partially across one of the two conductive layers.

A memory device having an array of storage cells is coupled to a third conductor extending partially across one of the two conductive layers. A data bus is coupled between the memory controller and the memory device for transferring data at a rate greater than the first clock rate and substantially twice the second clock rate (i.e., cycle frequency). Thus, if the memory device is DDR DRAM, then data is transferred across the bus at twice the second clock rate. However, the second clock rate is less than the first clock rate which is fed to the execution engine. If the memory device were an SDR SDRAM, for example, then a DDR DRAM will transfer data faster than SDR SDRAM even though the second clock rate is throttled downward compared to the first clock rate.

According to yet another embodiment, a method is provided for transferring data between a memory controller and a memory device. The method includes integrating an execution engine with a memory controller, and clocking the execution engine with a first clock that cycles at a first clock rate while clocking the memory controller with a second clock that cycles at a second clock rate. A memory device formed separate and apart from the integrated execution engine and memory controller is clocked with the second clock. However, data is transferred between the memory controller and the memory device on both the leading and trailing edges of the second clock.

30



The execution engine and memory controller are packaged upon a lead frame, with bonding pads coupled by a wire to leads extending from the package. Moreover, the memory device is also packaged in a lead frame with a wire connecting bonding pads to leads extending from the package. The packaged execution engine and memory controller are placed upon a printed circuit board (PCB) along with the packaged memory device. The PCB contains no more than two conductive layers spaced from each other by a dielectric layer. At least one of the two conductive layers is on an outer surface of the PCB, and contains bonding pads which receive the surface-mounted packaged integrated circuits.

### **BRIEF DESCRIPTION OF THE DRAWINGS**

Other objects and advantages of the invention will become apparent upon reading the following detailed description and upon reference to the accompanying drawings in which:

Fig. 1 is a block diagram of an execution engine and memory controller placed on a single monolithic semiconductor substrate, with the memory controller configured to control one or more arrays of semiconductor memory;

Fig. 2 is a block diagram of the execution engine and memory controller clocked at the same rate for transferring data synchronously between the execution unit and SDRAM semiconductor memory;

Fig. 3 is a block diagram of the execution engine clocked at a first rate and the memory controller clocked at a second rate slower than the first rate to transfer data to and from the DDR DRAM at twice the second rate; and

Fig. 4 is a state diagram for setting the clock to the memory controller at the second rate when power is applied to the memory controller;

Fig. 5 is a timing diagram of the first and second clock rates, and the effect on data transfer to and from the DDR DRAM when the clock rate is reduced from the first rate to the second rate;

5            Fig. 6 is a block diagram of DDR DRAM devices arranged on a surface of a PCB, along with a DDR DRAM controller, with trace conductors extending along two dielectrically spaced planes across the PCB;

10           Fig. 7 is a cross-sectional diagram of the PCB and trace conductors along plane 7 of Fig. 6; and

Fig. 8 is a circuit diagram of the circuit elements which drive and receive low power signals across the trace conductors.

15           While the invention is susceptible to various modifications and alternative forms, specific embodiments thereof are shown by way of example in the drawings and will herein be described in detail. It should be understood, however, that the drawings and detailed description thereto are not intended to limit the invention to the particular form disclosed, but on the contrary, the intention is to cover all modifications, equivalents and  
20 alternatives falling within the spirit and scope of the present invention as defined by the appended claims.

## **DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS**

Turning now to the drawings, Fig. 1 illustrates a memory module having a series of semiconductor memory devices 12 connected to a memory controller 16. A memory  
5 bus 11 is used to convey information between the memory devices 12 and a memory controller 16. Each semiconductor memory device 12 is preferably an asynchronous DDR dynamic random access memory (DRAM) device. Also, each memory device 12 includes an array of storage cells formed on separate chips of a semiconductor wafer, wherein the separate chips are removed and placed on a paddle of a lead frame. The  
10 paddle, chip, and a portion of the lead frame are preferably sealed inside a package. The package includes leads that extend from an outer surface of the package. Thereafter, the leads can be forwarded into holes or affixed to conductors on the surface of a PCB. There are various types of packages, some of which are through-hole dual in-line packages (DIP), j-leaded chip carriers, leadless chip carriers, etc., all of which have leads  
15 which extend preferably from an outer lateral surface of the package. An example of such packaging includes a small outline package (SOP).

There are numerous types of packaging technologies available in the marketplace. However, it is preferred that the chosen package be of very low cost. Therefore, SOP is  
20 preferably utilized as the package of choice. The packaging material can include multiple types of insulating substrates, such as ceramics or polyamides. The leads extending as part of a lead frame include any material which is electrically and thermally conductive, such as metals. Popular lead frame materials include copper, aluminum, tungsten, molybdenum, or a combination of the above.

25 Each semiconductor memory device 12 can be readily connected through known solder reflow techniques to bond the lower surface of the outward-extending lead to a corresponding upper surface of a bonding pad arranged on the upper surface of the PCB in registry with the corresponding lead. The bonding pads form one end of a printed  
30 conductor extending along a conductive layer of the PCB. The printed conductor can be formed by, for example, using a photosensitive material to chemically expose and remove

various portions of a planar sheet from the upper surface, leaving a pattern plurality of trace conductors. The upper layer of the trace conductors is formed on the board material, preferably made of an epoxy-bonded fiberglass, sometimes referred to as a FR-4 board. The epoxy-bonded fiberglass is also sometimes referred to as woven glass fabric, wherein special purpose materials, such as polyamide can be included into the fabric to enhance certain characteristics of the laminate material.

Referring again to Fig. 1, semiconductor memory device 12 is shown external to integrated circuit 10. Both memory 12 and integrated circuit 10 are shown having several I/O pins, some of which are reserved for control, address, data, and clock pins. Within integrated circuit 10 are two major subsystems: an execution engine 14 and a memory controller 16. Execution engine 14 is illustrated in part. Preferably, execution engine 14 is a processor. A processor is known to incorporate two operational units: an integer unit and a floating-point unit. These units communicate data via a data cache 18 drawn from, for example, external memory 12. The integer and floating-point units are controlled by a single instruction stream supplied by instruction cache 20. The instruction stream can be drawn from external memory 12, for example. Instruction cache 20 can supply multiple instructions per fetch if, for example, the processor is a superscalar processor. The actual number of instructions fetched can vary. As the instructions are dispatched from instruction cache 20, decoder 22 decodes those instructions and places them into prefetch unit 24.

Prefetch unit 24 works in conjunction with register file 26 and reorder buffer 28 to handle both in-order and out-of-order issues into functional units of execution unit 30. While it is generally easy to provide the required instruction bandwidth for sequential instructions since prefetch unit 24 can simply fetch several instructions per cycle in blocks of multiple instructions, it is much more difficult to provide instruction bandwidth in the presence of non-sequential fetches caused by, for example, branches.

Reorder buffer 28 works with the branch prediction techniques to align and merge back into the sequence of instructions various out-of-order instruction issues. The branch prediction can occur either in hardware or software by essentially predicting the outcomes of branches during instruction fetching without waiting for the functional unit 32 to indicate whether or not the branches should be taken. Prefetch unit 24 functions to hold all un-issued instructions regardless of the functional units (e.g., branch functional unit, ALU functional unit, shift functional unit, load functional unit, and store functional unit) within execution unit 30. Only the integer functional units are shown, prefetch unit 24 may be duplicated in the floating-point unit, depending on the implementation of the floating-point unit.

In addition to the functional units within the execution unit 30 producing various data that is temporarily stored in data cache 18, data can also arise from other functional units 32, external to the data flow path of the processor. Functional unit 32 can be attributed to, for example, the data flow path of an external data bus, such as a PCI bus. In addition to external functional units 32, data can also be input directly into the data flow stream or output from the data flow stream directly. The direct input and output of data is shown by block 34. For example, direct memory access can be implemented from a peripheral to and from semiconductor memory 12 without involving an instruction cycle of the processor.

The various blocks shown as reference numerals 18-34 can be synchronized to a system clock, where data is latched on either the rising or falling edges of the system clock. Memory controller 16, however, is synchronized to a controller clock. If it is desired that memory accesses by the execution engine be sequential and synchronized with the system clock, then the controller clock will transition at the same rate as the system clock. Instead, it is desired that memory accesses by the execution engine be asynchronous to accesses by the memory controller 16. Thus, memory controller 16 receives a controller clock that transitions at a different rate than the transitions of the system clock. It is the controller clock cycle that defines the bi-directional data strobe frequency on memory bus 11. Fig. 1 illustrates only the basic signals needed between

memory controller 16 and external memory 12. If DDR DRAM continues its price decrease, then it will be beneficial to implement DDR DRAM for memory device 12, rather than SDR SDRAM.

5            Fig. 2 illustrates the rudimentary features of a SDR SDRAM interface. If semiconductor memory 12 is SDRAM, then memory controller 16 is an SDRAM controller. SDRAM memory controller 16 is shown clocked by a controller clock that transitions at the same rate as the system clock sent to the execution engine 14. If the SDRAM is instituted as SDR, then the data transfer rate between controller 16 and  
10   SDRAM 12 is equal to the system clock cycle, or X MHz. Transferring data to and from the external memory is, therefore, synchronous with the operation of the execution engine and the memory controller (i.e., each subsystem utilizes the same clock cycle and clock frequency).

15            As DDR DRAM continues its price decrease, it is beneficial that the SDR memory device be replaced by DDR memory. Fig. 3 illustrates the various hardware elements used to perform the DDR interface and, more specifically, hardware elements that reduce the controller clock frequency relative to the system clock frequency. By controlling the memory controller with a lessened controller clock frequency, less  
20   complex packaging techniques can be used that are not sensitive to large parasitics while still utilizing the lower price-point DDR external memory devices. For example, during boot-up of integrated circuit 10, the basic I/O sequence (BIOS) of the computing system fetches bits of information stored within configuration register 40. If the proper bits are read, latch 42 will sample those bits and maintain a selectivity signal sent to multiplexer  
25   44. Alternatively, the value can be read-in from a reserved pin on integrated circuit 10. Regardless of how the data is read into latch 42, whatever data is present during boot-up is latched onto the output of latch 42 and presented to multiplexer 44.

             Multiplexer 44 is shown having two inputs: either the system clock input or the  
30   system clock having a reduced frequency. Divider 46 performs this reduction in frequency. If, for example, the reserved bit within configuration register 40 is of the

proper logic value, latch 42 will produce a selectivity signal that will select a reduced clock frequency instead of the system clock frequency. The configuration register 40 can be programmed during BIOS operations, for example, to store the configuration bit needed for activating either the full system clock or a reduced system clock. Preferably,  
5 the latch will maintain its output until power is removed from integrated circuit 10 during a power-on reset operation.

However, while power is maintained, instead of a first clock signal transitioning at X MHz frequency, a second clock signal transitioning at X-Y MHz is sent to memory  
10 controller 16. The second clock is sent in both the true and complementary forms as CLK and #CLK as shown. In addition, controller 16 forwards a separate reference supply voltage ( $V_{TT}$ ). The reference supply voltage is used to apply power to the output driver and receiver circuits, as well as applying pull-up voltage to termination resistors placed on each end of trace conductors which bear the data (DQ) signals that are sent  
15 according to the JEDEC standard.

Specifically, the JEDEC standard for DDR transmissions calls out the stub-series terminated logic (SSTL) specification by which the various address and data bus pins are terminated at both controller 16 and external memory 12. The various features of  
20 SSTL\_2 are described in the JEDEC standard, JEDEC Solid State Technology Association, December, 2000 (herein incorporated by reference). One benefit of the JEDEC standard is that the drivers and receivers are powered at a substantially reduced voltage value. For example, the voltage value for the data lines (DQ) is typically referred to as  $V_{DDQ}$ , which is defined to be less 2.5 volts or less. This allows for a lower power  
25 controller and memory device. The termination voltage or  $V_{TT}$  is essentially a reference voltage that is approximately equal to 1/2 of  $V_{DDQ}$ , for reasons described in the JEDEC standard.

For purposes of the description set forth in Fig. 3, memory controller 16 produces  
30 not only the signals needed for SSTL compatibility, but also the true and complementary clock signals to form the data strobe (DQS) signals. DQS is generally a strobe signal that

transitions at twice the true or complementary clock rate and dictates the data transfer rate of the bits of data sent across the data bus. SDR does not utilize a data strobe since data transfer is synchronized to the singular clock, not DQS that is twice the transfer rate of the singular clock.

5

By utilizing a reduced frequency controller clock, memory controller 16 can output data at a reduced rate relative to a memory controller that does not implement a reduced clock frequency. However, memory controller 16 nonetheless maintains compatibility to the lower-cost DDR memory. The lower-speed data rate allows the integrated circuit 10 and semiconductor memory 12 to be packaged with leads extending from a lateral surface, even though such packages and lead arrangements involve a higher parasitic capacitance and inductance. Avoiding flip-chip packaging and the costs associated therewith is one beneficial feature hereof.

15 Another beneficial feature is that instead of using low voltage TTL (LVTTTL) as the signal interface, DDR uses SSTL\_2. LVTTTL generally requires a power supply of approximately 3.3 volts, whereas SSTL merely requires a power supply of 2.5 volts or lower. Memory controller 16 thereby operates at a lower voltage level to achieve consideration power savings. Still further, DDR is made synchronous to whatever rate DQS happens to be, instead of fixing the DQS rate to be synchronous with the system clock, or twice the rate of the system clock. The DDR DRAM can, therefore, be made to operate asynchronous with the system clock.

SDR generally employs a N-prefetch architecture whereas DDR employs a 2N-prefetch architecture, where the internal data bus is twice the width of the external bus. This allows the internal memory cell to pass data to the input/output buffers in pairs. With DDR, there is no output enable for read operations, but DDR does support a burst terminate command to quickly end a read in process. During a write operation, the DM signal is available to allow the masking of non-valid write data. Since DDR utilizes a differential pair for the system clock, there must also be a true clock and a complementary clock. Positive clock edge for DDR refers to the point where the rising

30



clock signal crosses with the falling complementary clock signal, and the term negative clock edge indicates the transition of the falling clock and the rising complementary clock signals.

5            Fig. 4 illustrates a state diagram of the set and reset states of latch 42 (Fig. 3). A BIOS operation 50 can occur during a power-on reset (PWR\_ON) at which the computing system is initiated. During a set state 52, the voltage values within the configuration register are read (RD\_CONFIG\_REG). Set state 52 remains in the set condition to latch whatever is present on the set input of the latch until reset state 54  
10        occurs. The reset state can occur during a power-off (PWR\_OFF). Once the latch is reset, it remains reset until power-on occurs or, in the alternative, a set state is manually initiated.

             Fig. 5 illustrates a timing diagram of various signals sent between the memory  
15        controller and the external memory device. More specifically, Fig. 5 shows the effects of reducing the memory controller clock on those signals. For example, the timing diagram to the left of line 56 indicates signals transitioning as a result of the controller clock transitioning at a first clock rate of X MHz. The timing diagram to the right of line 56  
20        indicates what occurs when the controller clock is throttled to a lower frequency of X-Y MHz. As shown, the true and complementary clocks sent from the memory controller to the memory devices slows down as the control clock signal is throttled down. Moreover, the time slots for controlling and addressing the data is increased, along with the data  
25        strobe signal. Since data is captured on both the rising and falling edges of the data strobe signal, the data transfer rate which follows the second clock signal is reduced relative to the data transfer rate that follows the first clock signal (i.e., the first clock  
30        signal and its effects are to the left of line 56 and the second clock signal and its effects are to the right of line 56).

             Fig. 6 illustrates a plan view of a PCB 60 having various packaged chips placed  
30        on an upper surface of PCB 60. Packaged chips 62a and 62b can be packaged semiconductor memory 12 (Fig. 3) and, specifically, DDR DRAM. Also placed on PCB

60 is a packaged integrated circuit 64, which is preferably an execution engine integrated with a memory controller. Alternatively, packaged integrated circuit 64 can simply be a memory controller. Various trace conductors 66 are shown extending between, for example, edge connectors 68 and the leads of packaged chips 62 and 64.

5

Edge connectors 68 can be used to slidably engage into a receptacle, for example. Of course, there could be numerous trace conductors that extend exclusively between packaged chips, and do not extend from a packaged chip to edge connector 68. In addition to transferring signals, trace conductors 66 can also transfer power and ground supplies between chips and/or between edge connector 68 and one or more chips. Trace conductors 66 are made from a conductive material as described above, and are arranged coplanar to each other on two conductive layers separated by a dielectric layer. The power and signal-bearing conductors can be formed on both conductive layers.

Fig. 7 illustrates a cross-sectional view along plane 7-7 of Fig. 6. Specifically, Fig. 7 illustrates PCB 66 having a maximum of two conductive layers. One conductive layer 70 can be dielectrically spaced below another conductive layer 72. Each conductive layer can contain trace conductors for conveying transitory signals, and steady-state power and ground. Preferably, whatever space is not occupied by a signal trace conductor will be consumed with a power and ground trace conductor. Thus, the power and ground trace conductors can be considerably wider than the signal trace conductors. As shown, power and ground trace conductors 74 can be arranged on the upper and/or lower conductive layers to fill the unused areas not bearing signal trace conductors 76. Also shown is a via 78 which might connect a signal trace conductor 76 on an upper surface to a signal trace conductor 76 on the lower conductive layer.

Contrary to PCBs having four or more conductive layers, with possibly separate conductive layers reserved for a power supply plane and a ground supply plane, using only two conductive layers is considerably less expensive to manufacture than PCBs having four or more conductive layers. The degree of complexity further increases whenever vias must extend across two or more intermediate conductive layers. Using a

maximum of two conductive layers substantially reduces the cost and complexity of the overall design.

Fig. 7 also illustrates the interior configuration of packaged chips 62b and 64. As shown, the chip 80 can be placed on a paddle 82 within a lead frame. A bonding pad on the upper surface of chip 80 can, therefore, receive a wire bond 84 which extends from the bonding pad to a bond finger also associated with the lead frame. The bonding finger extends outward as a contiguous conductive member from the outer surface of the ceramic or polyamide package, as a lead 86 that is then solder bonded to a corresponding trace conductor 76. Although not shown but understood, adhesives may be used to secure the underside of packaged chips 62b and 64 to the upper surface of PCB 60. The adhesive is preferably thermally insulated and substantially non-conductive.

Fig. 8 illustrates the preferred mechanism by which each trace conductor between the external memory device and the memory controller, as well as between memory devices, is carried out. Specifically, Fig. 8 illustrates SSTL which includes SSTL versions 1 or 2, oftentimes referred to as SSTL and SSTL\_2. At the output of the controller or memory device is a driver 90. The output driver can be achieved by connecting one or a pair of inverters together between a power supply and ground. The output of the inverter is then fed into the trace conductor which has some series resistance ( $R_S$ ). The trace conductor then transfers the corresponding signal to the receiver 92. Indigenous to SSTL technology is a pull-up resistor ( $R_T$ ). The pull-up resistor is used to pull the signal up to the termination voltage ( $V_{TT}$ ). Preferably, the termination voltage is equal to a reference voltage ( $V_{REF}$ ). In some applications, the reference voltage is approximately one-half the device supply voltage ( $V_{DD}$ ). The device supply voltage is also used as a supply voltage for the memory device.

According to one example, the receiver functions as a comparator that compares the input voltage ( $V_{IN}$ ) to the reference voltage ( $V_{REF}$ ). If  $V_{IN}$  exceeds  $V_{REF}$ , then an output will be produced according to well-known comparator and differential pair common source amplifier technology. Because of the high-speed nature at which signals

are transferred in DDR memories and the use of SSTL termination at each end of the corresponding trace conductors, the memory controller can maintain relatively high speed data transfers yet at a reduced power supply and with lessened power consumption.

5           Although the invention has been described in connection with specific and preferred embodiments, it should be understood that the invention as claimed should not be unduly limited to such specific embodiments. Specifically, the memory controller is claimed as being either integrated along with the execution unit or separate from the execution unit (i.e., as a single chip or as two chips). Moreover, the PCB has a maximum  
10 of two conductive layers, wherein at least one conductive layer is on an external surface. Both conductive layers can be on the opposed external surfaces if desired. Preferably, however, one conductive layer is embedded between the external surfaces.

          Preferably, the memory controller and external memory uses an asynchronous  
15 clock design, whereby the system clock frequency is different from the memory controller and external memory clock frequencies. Simply as an example, the system clock frequency might be at 150 MHz whereas the memory controller clock frequency might be reduced to 100 MHz. Since DDR is employed, the data rate is at 200 MHz. By operating the I/O interface on the memory controller at less than the system clock speed  
20 which toggles the execution unit and less than the rated speed of the DDR external memory, the raw data rate across the DDR memory data bus will be higher than the data rate achieved with SDR SDRAM, for example. However, since the clock speed of the I/O interface is lower than that used with SDR SDRAM, the power consumption within the integrated circuit which bears the memory controller will be lower using a DDR  
25 interface than when using an SDR SDRAM interface. Even though the external memory might be rated at speeds exceeding 266 MHz, for example, the external memory will only be clocked at 200 MHz (assuming in this example that the I/O interface is clocked by a control clock transitioning at 100 MHz).

30

Performance enhancement is achievable through extensive use of burst-mode transfers and reliance on on-chip cache memory. For example, 100 MHz DDR DRAM clock can be used instead of, for example, 150 MHz SDRAM clock. The reduced clock signal reduces power consumption, as well as parasitic capacitance and inductance at the memory controller pin-out. This will allow use of conventional lead frame packaging and two-layer PCB rather than the more expensive flip-chip packaging and four (or more) layer PCBs. In addition, less expensive DDR DRAM can be used rather than SDR SDRAM, for example. Still further, DDR DRAM manufacturers can specially market lower frequency external memory devices since the present interface does not require operation at the rated speed of that external memory. Thus, DDR DRAMs which might pass 200 MHz functionality but not 266 MHz can be marked (and marketed) to a lower frequency design and, therefore, higher yields might be achievable at the lower speed functional test.

Implementing lower rated speeds and cheaper packaging techniques, as well as cheaper PCBs, provides significant cost savings with enhanced marketability for certain price points. Accordingly, the structures and methods described herein are within the scope and spirit of the present claims, and their equivalents are thereby covered as set forth in the appended claims.